THE TIME-SERIES PROPERTIES OF AGGREGATE CONSUMPTION: IMPLICATIONS FOR THE COSTS OF FLUCTUATIONS

Ricardo Reis
Columbia University

Abstract
The properties of the stochastic process followed by aggregate consumption affect the estimates of the costs of fluctuations. This paper pursues two approaches to modeling consumption dynamics and measuring how much society dislikes fluctuations, one statistical and one economic. The statistical approach estimates the properties of consumption and calculates the costs of having consumption fluctuating around its mean growth. The paper finds that persistence is a crucial determinant of the costs and that the high persistence in the data severely distorts conventional measures. It shows how to compute valid estimates and confidence intervals. The economic approach uses a calibrated model of optimal consumption and measures the costs of eliminating income shocks. This uncovers a further cost of uncertainty, through its impact on precautionary savings and investment. The two approaches lead to costs of fluctuations that are higher than the common wisdom, between 0.5% and 5% of per capita consumption. (JEL: E32, E21, E60)

1. Introduction

In a famous contribution, Robert Lucas Jr. (1987) asked: What would be the effect on welfare of eliminating economic fluctuations? As Lucas (page 3) put it, answering this question would allow us "to get a quantitative idea of the importance of stabilization policy relative to other economic questions." To reach an answer, Lucas made three assumptions. First, he assumed that society's preferences can be represented by a welfare function that depends only on the time path of consumption per capita. That is, he assumed not only that there is a representative consumer, but also that her utility function represents society's normative preferences. Second, he assumed that this welfare function is time-separable and iso-elastic. Third, he assumed that the log of annual per capita consumption is

The editor in charge of this paper was Jordi Galí.

Acknowledgments: I am grateful to Gadi Barlevy, Jordi Galí, Per Krusell, Albert Marcet, Yves Nosbusch, Jonathan Parker, Chris Sims, Lars Svensson, Mark Watson, and anonymous referees for useful comments.

E-mail address: rreis@columbia.edu

Journal of the European Economic Association June 2009 7(4):722–753

serially uncorrelated and normally distributed around a linear trend. These three assumptions produced a surprising result: Society would be willing to sacrifice a meager 0.05% of consumption to get rid of fluctuations. The economic fluctuations that macroeconomists have focused so much attention on, cost each person on average only $12 per year.

A large literature has followed focusing especially on the first two assumptions. Imrohoroğlu (1989), Atkeson and Phelan (1994), Krusell and Smith (1999), Storesletten, Telmer, and Yaron (2001), Beaudry and Pages (2001), and Krebs (2003, 2007) measured the costs of fluctuations in economies where agents are heterogeneous and markets are incomplete, so that there is not a representative consumer whose preferences are a valid measure of welfare. Although it is conceivable that the costs of fluctuations would be higher, as bad income shocks hurt a few households severely, the typical finding from these studies is that the costs of fluctuations are only slightly higher or even lower than the Lucas benchmark. Other studies looked at the second assumption of iso-elastic preferences. Dolmas (1998), Otrok (2001), Tallarini (2000), and Epaulard and Pommeret (2003) assumed different utility functions, whereas Alvarez and Jermann (2004) used asset prices to elicit rather than assume preferences over risk. Although many of these studies found larger estimates of the costs of fluctuations, this typically came at the expense of assuming that people are extremely averse to risk, which appears to be inconsistent with the risk-taking that we observe in their choices (Lucas 2003).

The focus of this paper is on the third assumption that consumption is serially uncorrelated. I will present alternative models of consumption dynamics and study their impact on estimates of the costs of fluctuations. First, I will consider statistical models of aggregate consumption and show that if consumption is very persistent, as is the case in the U.S. data, Lucas's (1987) estimates are severely downward-biased. A methodological contribution of this paper is to show how to construct reliable estimates of the costs of fluctuations using a small sample of data when there is persistence of the degree that we observe. Second, I will consider economic models in which consumption fluctuations are an optimal response to shocks. One virtue of having endogenous consumption choices is that it uncovers a further cost of fluctuations through its effect on the level and growth rate of consumption via the desire for precautionary savings and risky investments. The discipline imposed across the different models is that they must all match the main features of the consumption data.[1]

The Lucas assumption that shocks to consumption are serially uncorrelated is clearly dismissed by the data. More surprisingly, using either type of model,

1. The "statistical models" can be interpreted as economic models of endowment economies. The distinction between economic and statistical models in this paper is simply between those where the path of consumption is chosen or not, respectively.

the adequate process for aggregate consumption implies that the costs of fluctuations are actually one or two orders of magnitude larger than what Lucas argued: between 0.5% and 5% of per capita consumption.

It is important to be clear from the outset how these estimates should be used. This paper does not use "fluctuations" and "business cycles" interchangeably. Its focus is on the former, and the aim is to measure the costs of eliminating the uncertainty that makes consumption fluctuate. These estimates do not distinguish between fluctuations due to productivity or monetary shocks, or between those that correspond to business cycles and those that are due to uncertainty about long-run growth. Likewise, it does not distinguish between efficient or inefficient fluctuations, or between those that policy can do something or nothing about. What these large numbers suggest is that focusing attention on deterministic growth models, as happened at least partly in response to Lucas's original results, will be missing out on a significant part of welfare. Section 6 will discuss how to use the results to assess the costs of business cycles and the role for policy.

The contributions of this paper are: (i) to show that the time-series properties of the model of consumption and especially its persistence are important inputs into the costs of fluctuations, (ii) to propose different strategies to estimate and model it, and as a result (iii) to provide more reliable numbers for the costs of fluctuations. This paper is a member of a recent family of work that highlights the role of the statistical properties of consumption in the costs of fluctuations. Weitzman (2007) emphasized the role of thick-tails in the distribution, and Barro (2006) the effect of rare large shocks.

Whereas the assumption that consumption is serially uncorrelated has received little attention in the literature, a few papers assumed instead that consumption follows a random walk (Dolmas 1998; Tallarini 2000; Epaulard and Pommeret 2003). Their focus, however, was on the other assumptions behind the Lucas calculations. This paper investigates more systematically the time-series properties of consumption, and its persistence more specifically. It goes beyond asking whether consumption is stationary or not, but instead tries to obtain good measures of its persistence and their impact on the costs of fluctuations. This focus leads the paper to address problems with making inferences about the costs of fluctuations using small samples that the literature has so far ignored.[2]

More related to this paper is Obstfeld (1994), who found that, for reasonable calibrations, the costs of fluctuations are small even if consumption is infinitely persistent. This paper reaches the opposite conclusion and shows that the difference is due to the way a key parameter is calibrated: the effective discount rate.

2. Pallage and Robe (2003) also look at the stochastic properties of consumption and the costs of fluctuations, but to compare them across many developing countries.

Section 2 discusses the role of this parameter on the costs of fluctuations and argues that calibrating it to fit the data reverses Obstfeld's conclusion.[3]

Tallarini (2000) and Otrok (2001) used more elaborate economic models of consumption designed to better fit the properties of business cycles. Tallarini emphasized the implications of risk-sensitive preferences for the costs of fluctuations, whereas Otrok's estimates do not capture the precautionary-investment effect. The economic models in this paper are simpler, and likely fit the data worse, but they allow me to focus on the two key features that this paper emphasizes: persistence and precautionary-investment effects. Like Barlevy (2004), this paper measures the costs of fluctuations without excluding the possibility that these may have long-lasting effects, either through long-lived fluctuations or through an impact on the average growth rate.

Finally, Alvarez and Jermann (2004) also concluded that the costs of fluctuations may be large because of persistent shocks, but by following a very different approach. They used asset pricing data to infer the marginal utility of consumption (and implicitly risk aversion), tackled the problem of accurately estimating risk premia and covariances between consumption and asset prices, and emphasized the need for a model of how consumers trade risk. This paper instead assigns low values to risk aversion but uses asset price data to infer the effective discount rate, tackles the problem of accurately measuring the persistence of consumption, and emphasizes the need for a model of consumption dynamics over time. In some ways, it is remarkable that we reach similar conclusions.

The paper proceeds as follows. Section 2 presents some simple models of consumption that highlight the main determinants of the costs of fluctuations. These involve choosing one key parameter, the effective discount rate, and Section 3 discusses how to pick its value. Section 4 estimates the costs of fluctuations across a variety of statistical models for consumption, and Section 5 uses instead economic models. Section 6 concludes by interpreting the economic significance of the estimates.

2. Models of Consumption and the Costs of Fluctuations

The starting point for most studies of choice under uncertainty is that people dislike risk. Faced with a choice between its current risky consumption series $\{C_t\}$ and a "suitably modified" consumption series $\{\bar{C}_t\}$ that is purged from fluctuations, it is assumed that society would choose the latter. Lucas (1987) emphasized that one can go one step further and quantify this preference for stability. He suggested measuring the costs of fluctuations by the fraction of annual consumption

3. Van Wincoop (1999) makes a related point in the study of the benefits from international risk-sharing.

that society would be willing to pay to eliminate these fluctuations. Maintaining his assumptions of a utility function that is time-separable (with subjective discount rate β) and iso-elastic (with a coefficient of relative risk aversion γ), the costs of fluctuations are defined as the scalar λ that solves the equation

$$E\left[\sum_{t=0}^{\infty} e^{-\beta t} u(C_t(1+\lambda))\right] = \sum_{t=0}^{\infty} e^{-\beta t} u(\bar{C}_t), \tag{1}$$

where $u(C_t) = \ln(C_t)$ if $\gamma = 1$ and $C_t^{1-\gamma}/(1-\gamma)$ otherwise, and $E[\cdot]$ denotes the expectation operator conditional on information at time 0.

Solving this equation requires two pieces of information: the stochastic process for the risky consumption series $\{C_t\}$ in order to evaluate the expectation on the left-hand side, and the precise definition of the unobservable counterfactual consumption series $\{\bar{C}_t\}$. Both of these requirements are met by having a model for consumption. This paper will consider two distinct approaches to modeling consumption: One consists of estimating a statistical process for consumption; the other consists of assuming an economic environment in which society optimally chooses how much to consume.

2.1. Statistical Models of Consumption

From a statistical perspective, a natural choice for the counterfactual consumption series is expected consumption. Eliminating fluctuations then corresponds to eliminating the variability of consumption, while keeping its mean unchanged. In the U.S. economy in the past century, consumption has grown at an approximately constant rate g. An appropriate model for counterfactual consumption is $\bar{C}_t = E[C_t] = C_0 e^{gt}$.

I will maintain the assumption that consumption is log-normally distributed, because it is consistent with the U.S. data and it is analytically convenient.[4] Appendix A.1 shows that the costs of fluctuations are

$$\ln(1+\lambda) = \begin{cases} 0.5(1-e^{-\rho}) \sum_{t=0}^{\infty} e^{-\rho t} \operatorname{Var}(c_t) & \text{if } \gamma = 1, \\ (\gamma-1)^{-1} \ln\left[(1-e^{-\rho}) \sum_{t=0}^{\infty} e^{-\rho t} e^{0.5\gamma(\gamma-1)\operatorname{Var}(c_t)}\right] & \text{if } \gamma \neq 1, \end{cases} \tag{2}$$

where $c_t = \ln(C_t)$ and $\rho = \beta + (\gamma - 1)g$ is the effective discount rate that weighs future costs.

The costs of fluctuations with log-normality depend solely on one property of consumption: its forecast error variance at different horizons. This is determined

4. However, see Weitzman (2007) for a discussion of the possibility that log consumption instead has a t-distribution.

by the dynamics of the stochastic component of log consumption, $\hat{c}_t$, and a simple model for it is

$$\hat{c}_t = \eta \hat{c}_{t-1} + \varepsilon_t, \tag{3}$$

where ε_t is normally distributed with mean zero and variance σ^2. This representation fits the post-war U.S. consumption data well: Lagged consumption accounts for 84% of the variability of present consumption when η equals the least squares estimate 0.92. Moreover, special cases of equation (3) correspond to two important processes. Lucas (1987) assumed that $\eta = 0$ and I will correspondingly call this the *Lucas consumption process*. Hall (1978) showed that rational expectations approximately predict that $\hat{c}_t$ follows a random walk and that the U.S. data is consistent with this assumption. This corresponds to $\eta = 1$, which I will label the *Hall consumption process*.

With this AR(1) model and if $|\eta| \leq 1$, Appendix A.2 shows that the costs of fluctuations approximately equal

$$\lambda \cong \frac{0.5\gamma\sigma^2}{\rho + 1 - \eta^2} \tag{4}$$

$$= \frac{0.5\gamma(1-\eta^2)}{\rho + 1 - \eta^2} \times \left(\frac{\sigma^2}{1-\eta^2}\right). \tag{5}$$

The costs of fluctuations therefore depend on the value of four parameters. Two of these regard preferences, γ and ρ, and will be discussed in Section 3. The focus of this paper is on the other two, σ^2 and η, that capture the properties of the stochastic process for consumption.[5]

The first expression (4) shows that λ increases with both the variability and the persistence of consumption. The larger is the variability of shocks to consumption, the more society finds these shocks costly, so the more it is willing to pay to eliminate consumption fluctuations. The more persistent are shocks to consumption, the more long-lived is their impact on consumption, and thus the larger their cost. Still, for $\rho = 0.02$, which Section 3 will justify, even when η is as high as 0.8 so that a shock to consumption takes about two years to dissipate by half, the costs of fluctuations are only twice higher than those with a process with no persistence. As persistence increases further though, the costs of fluctuations increase quite rapidly. If η is 0.9, the costs are already 7 times larger than with a Lucas process, and if $\eta = 0.95$ they are 14 times higher. Even for stationary processes, high persistence can significantly raise the costs of fluctuations. The effect

5. When $\eta = 0$, the formula in equation (4) differs from the one derived by Lucas (1987) by a factor of $1/(1+\rho)$. This difference arises because I evaluate expected utility conditional on information at time 0, whereas Lucas computes the unconditional expectation. Because ρ is close to zero, this difference is quantitatively negligible. I focus on the conditional rather than the unconditional expectations, since in the latter case the costs of fluctuations would be infinite when $\eta = 1$ and would be severely downward biased when η is close to 1 since the unconditional variance would be estimated using the relatively short post-war U.S. sample.

is more dramatic when we shift from the Lucas to the Hall models. If $\rho = 0.02$, then the Hall consumption model predicts costs of fluctuations that are 51 times larger than those estimated by Lucas and if $\rho = 0.01$, the costs of fluctuations are two orders of magnitude larger than what Lucas estimated.

These calculations assumed that σ^2 was held fixed while η varied. It might be argued that Lucas (1987) instead measured the unconditional variance of consumption, which corresponds to $\sigma^2/(1-\eta^2)$. In expression (5), the first term actually decreases as η rises. The reason is that keeping the unconditional variance fixed, raising η increases the predictability of consumption by lowering its forecast error variance. The consumer therefore faces less risk so the costs of fluctuations fall. Rather than undermining the argument of the previous paragraph, this alternative view of the Lucas calculation provides another way to see its limitations. Lucas used a finite sample to gauge the unconditional variance of consumption. This implies that if consumption is very persistent, his estimate is severely downward-biased. This is particularly clear in the case where consumption follows a random walk: while in a finite sample one obtains a finite estimate of the variance of consumption, the actual variance is infinite. Even if consumption is stationary, if it is very persistent, one will obtain a very downward-biased estimate of its variance using the post-war U.S. sample.[6]

Whichever way you look at it, these calculations show that it is crucial to jointly estimate both the volatility of shocks to consumption and their persistence. Section 4 will attack this estimation problem directly using different statistical approaches.

2.2. Economic Models of Consumption

An economic model is a specification of the environment facing a representative consumer earning a random income stream, and the optimal consumption choices are the process to be considered.[7] The counterfactual consumption with no fluctuations is what the consumer would choose if income were stable.

In this section, I consider a simple economic environment. The consumer solves

$$\max_{\{C_t\}} \quad E\left[\sum_{t=0}^{\infty} e^{-\beta t} u(C_t)\right] \tag{6}$$

$$\text{subject to:} \quad K_{t+1} + C_t = R_t K_t. \tag{7}$$

6. The downward bias in the estimate of $\sigma^2/(1-\eta^2)$ follows directly from the downward bias in η. Andrews (1991, Section 8) reports Monte Carlo experiments with very large biases in the estimation of the unconditional variance even for η as low as 0.9.

7. To focus solely on the third of the Lucas (1987) assumptions, I maintain the assumption of a representative consumer. It would be interesting in future work to take into account the large idiosyncratic risks facing households (Parker and Preston 2004), while modeling their consumption processes carefully.

The budget constraint states that savings (K_{t+1}) plus consumption equals income. Last period's savings are the only source of income through investment in a risky technology with positive marginal return R_t, which is log-normally distributed with mean $r - 0.5\sigma^2$ and variance σ^2. The consumer starts at time 0 with some positive amount of capital K_0.

Appendix A.3 shows that the solution to this problem is

$$c_t = c_{t-1} + g - 0.5\sigma^2 + \varepsilon_t, \tag{8}$$

where

$$g = r - \rho + 0.5\gamma(\gamma + 1)\sigma^2 - \gamma\sigma^2, \tag{9}$$

with initial condition

$$C_0 = (1 - e^{g-r})R_0 K_0. \tag{10}$$

Consumption is log-normal with expected growth g, and its log follows a random walk as in equation (3) with $\eta = 1$.[8]

However, there is one important difference between this economic model and its statistical counterpart. In the economic model, both the level C_0 and the growth rate of consumption g are functions of σ^2. Income uncertainty not only causes fluctuations in consumption but also has two effects on the level and growth rate of consumption, captured by the two terms on the right-hand side of the expression for g. The first effect is due to precautionary savings: The rational consumer reacts to the uncertainty by saving more. This allows her to accumulate a stock of precautionary savings to safeguard against unexpected future bad shocks. The second effect is due to investment risk: The risk-averse consumer will shy away from investing in the risky technology. In this model, as long as relative risk aversion exceeds one, the combined precautionary-investment effect is such that eliminating fluctuations would raise the level of consumption and reduce growth.[9]

The counterfactual $\bar{C}$ is defined by equations (8)–(10) with $\sigma^2 = 0$. It therefore differs from average consumption both in the level of initial consumption and in its growth rate. Although one can follow Lucas and calculate the gains from eliminating fluctuations in consumption, one needs a theory of consumption choices to calculate the costs of fluctuations in income. The latter affect not just the fluctuations in consumption, but also the level and growth rate of consumption through the precautionary-investment motive. Although this model fits

8. There is a $0.5\sigma^2$ term in equation (8) to ensure that $E_0(C_t) = C_0 e^{gt}$.

9. Angeletos (2007) provides a thorough study of these two effects and a general characterization of the conditions for one to dominate the other. Barlevy (2004) has suggested a complementary channel through which fluctuations affect growth. Eliminating uncertainty may raise investment in innovative activities and consequently long-run growth.

into the statistical model (3) with $\eta = 1$, Appendix A.3 shows that the costs of fluctuations are now higher (if $\gamma > 1$) because of a new term in the denominator:

$$\lambda \cong \frac{0.5\gamma\sigma^2}{\rho - 0.5\gamma(\gamma - 1)\sigma^2}. \tag{11}$$

Moreover, note that this precautionary-investment effect is more general than the model in this section. It will be present in most economic models of consumption under uncertainty, regardless of their predictions for the persistence of consumption.[10] Likewise, although growth may be higher or lower without uncertainty, welfare will always be higher. By ignoring this effect, statistical models may underestimate the costs of fluctuations.

2.3. Initial Estimates of the Costs of Fluctuations

Table 1 presents estimates of the costs of fluctuations for the different models that I have discussed so far.[11] The value of σ^2 for each model is estimated using U.S. annual data from 1947 to 2003 on real per capita consumption of nondurables and services from the Bureau of Economic Analysis. This will be the measure of consumption used in this paper. Quarterly data leads to very similar results; total consumption, which inappropriately includes expenditure on durables as current consumption, approximately doubles the estimate of σ^2 and so doubles all of the estimates of the costs of fluctuations. The values for γ and ρ will be discussed in Section 3.

Panel (A) displays the estimates with the Lucas model of consumption. As Lucas (1987) originally concluded, fluctuations cost very little, between 0.04% and 0.2% of per capita consumption. Panel (B) presents estimates for the AR(1) statistical model fitted to the U.S. data. The estimated η implies a considerable amount of persistence, with a half-life of deviations from trend growth after a shock of 8 years. However, the estimated costs still lie in the same range as the Lucas estimates.[12] These results should be interpreted with caution though: It is well known that for very persistent processes, least-squares estimates are statistically inconsistent and severely downward biased.

Panel (C) shifts to the economic model presented in this section. The infinite persistence of shocks and the precautionary savings effect combine to generate

10. Epaulard and Pommeret (2003) find an effect of volatility on growth in an AK-growth model, but interpret it as being specific to endogenous growth models. Actually, this effect is present in most models of consumption and uncertainty.

11. These numbers use the exact formula in equation (1) rather than the approximations in equations (4) and (11), although the two are typically identical to the second decimal percentage point.

12. The reader may be surprised that the estimates in panel (B) are actually lower than those in panel (A), in spite of the higher persistence. The reason is that the estimated volatility of shocks is lower for the AR(1) than for the Lucas model, which drives down the costs of fluctuations.

TABLE 1. The costs of fluctuations in three simple models.

	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
Panel A: The Lucas statistical model			
	0.04%	0.12%	0.20%
	(\$9)	(\$28)	(\$46)
Panel B: The AR(1) statistical model estimated by least squares			
$\rho = 0.03$	0.03%	0.10%	0.17%
	(\$8)	(\$24)	(\$40)
$\rho = 0.02$	0.04%	0.11%	0.18%
	(\$8)	(\$25)	(\$43)
$\rho = 0.01$	0.04%	0.12%	0.19%
	(\$9)	(\$27)	(\$45)
Panel C: The random walk economic model			
$\rho = 0.03$	0.21%	0.62%	1.03%
	(\$48)	(\$145)	(\$242)
$\rho = 0.02$	0.31%	0.94%	1.56%
	(\$73)	(\$219)	(\$365)
$\rho = 0.01$	0.63%	1.88%	3.14%
	(\$147)	(\$441)	(\$735)

Notes: Each cell shows the per capita costs of fluctuations as a fraction of consumption and, in parentheses, in U.S. 2003 dollars. The standard deviation of shocks is 0.028, 0.011, and 0.011, for Panels (A) to (C), respectively.

substantially larger costs of fluctuations, between 0.2% and 3.1%. This upper bound is almost 80 times larger than the smallest number in Panel (A) that Lucas focused on.

After a brief detour in the next section to discuss the calibration of γ and especially ρ, the remainder of this paper explores more refined estimates of the costs of fluctuations. Section 4 estimates statistical models that deal with the high persistence of consumption data, and Section 5 builds more elaborate economic models of consumption.

3. Choice of Parameters and the Effective Discount Rate

As equations (4)–(5) show, the higher γ, the higher the costs of fluctuations, because the more people dislike risk. The choice of the coefficient of relative risk aversion has been extensively debated in the literature. The conventional choices for γ are between 1 and 5, and these are the numbers that I use.

As for ρ, it measures the effective rate at which consumers discount the impact of shocks on future consumption. The smaller it is, the less people discount the future costs of a persistent shock so the larger the overall costs of fluctuations. This section discusses the calibration of ρ.

3.1. Calibrating the Effective Discount Rate

Ramsey (1928) famously showed that without uncertainty $\gamma g = r - \beta$. From this and the definition of ρ, it follows that $\rho = r - g$, relating the effective discount rate to two observables: the average return on savings and the growth rate of consumption. With uncertainty, there is an extra term involving the variance of consumption, as we can see in equation (9). Quantitatively though, given how small σ^2 is in the data, this term is negligible, at most 0.1%.

Poterba (1998) estimates the after-tax return on capital in the United States in the period 1959–1996 to be either 3.9% or 5%, depending on whether one includes property taxes or not. McGrattan and Prescott (2003) use data from 1880 to 2002 and find returns of 4% on accounting capital, and 5.4% on equity. As for the average annual growth rate of per capita consumption, it equals 2.2%. These point estimates therefore suggest a value for ρ somewhere between 1.7% and 2.2%. Correspondingly, I will consider the values of 1%, 2%, and 3% for ρ.

3.2. The Link Between Effective and Subjective Discount Rates

It is preferable to calibrate the effective discount rate ρ, rather than the subjective discount rate β, for at least three reasons.

First, because ρ is observable from data on interest rates and consumption growth, while β is not observable.

Second, because ρ is the rate that matters. This is clear by staring at equation (2), but it can be seen more generally. The costs of fluctuations depend on $\hat{C}_t$, the stochastic part of consumption, which equals $C_t e^{-gt}$. With iso-elastic preferences, this implies that welfare equals $\sum_{t=0}^{\infty} e^{-\rho t} u(\hat{C}_t)$, where recall $\rho = \beta + (\gamma - 1)g$.[13] It is at rate ρ, not β, that fluctuations in consumption are discounted. Another way to see that the effective discount rate is the one that matters is to note that while for some pairs (γ, ρ), the implied β will be negative, welfare remains bounded from above as long as the effective discount rate is positive (Kocherlakota 1990).

Third, setting β can lead to incorrect inferences and counter-intuitive results. For instance, in the formula for the costs of fluctuations in equation (4), raising the coefficient of relative risk aversion raises the costs directly, but lowers them indirectly by raising ρ. It is easy to come up with choices of parameters for which fixing β implies that higher risk aversion leads to lower costs of fluctuations.[14] The intuition is that a higher γ implies a lower intertemporal elasticity of substitution

13. If $\gamma \neq 1$, then $\sum_{t=0}^{\infty} e^{-\beta t} u(C_t) = \sum_{t=0}^{\infty} e^{-\rho t} u(\hat{C}_t)$, whereas if $\gamma = 1$, the two sides differ by a constant.

14. I am grateful to Per Krusell for bringing this to my attention.

so that society discounts the future at a higher rate, thus lowering the weight of future uncertainty.

3.3. *The Role of the Intertemporal Elasticity of Substitution*

This last argument raises the role of intertemporal substitution, as distinguished from relative risk aversion, on the costs of fluctuations. As Obstfeld (1994) noted, risk aversion determines the per-period cost of volatility, whereas intertemporal substitution determines the weights given to the future cumulative per-period costs.

With the preferences in equation (8), the elasticity of intertemporal substitution equals the inverse of relative risk aversion, so the two concepts cannot be distinguished. To investigate further, consider the specification of preferences due to Epstein and Zin (1989) and Weil (1990). Utility at date t, V_t, is defined by the recursion

$$\begin{aligned}&[1+(1-e^{-\beta})(1-\gamma)V_t]^{(1-\theta)/(1-\gamma)}\\ &\quad=(1-e^{-\beta})C_t^{1-\theta}+e^{-\beta}[1+(1-e^{-\beta})(1-\gamma)E_t[V_{t+1}]]^{(1-\theta)/(1-\gamma)}. \qquad (12)\end{aligned}$$

The parameter γ still equals the coefficient of relative risk aversion, but now the intertemporal elasticity of substitution is $1/\theta$.[15]

Appendix A.4 solves for optimal consumption and for the costs of fluctuations in the economic model in equations (12) and (7) and shows the following surprising result: *The costs of fluctuations with Epstein–Zin–Weil preferences (12) are the same as the costs with iso-elastic preferences (6) up to a term in $O(\sigma^4)$*. Therefore, distinguishing between intertemporal substitution and risk aversion does not affect the estimates of the costs of fluctuations. Moreover, the intertemporal elasticity of substitution does not enter the formula for the costs of fluctuations.

How can this finding be reconciled with Obstfeld (1994)? With Esptein–Zin–Weil preferences, $\rho=\beta+(\theta-1)g$. Obstfeld calibrated β, so when he lowered the elasticity of intertemporal substitution by raising θ, he raised his value for ρ. The result in the previous paragraph instead keeps ρ fixed. This provides another illustration of why it is better to calibrate ρ rather than β. The counterpart expression to equation (9) with the preferences in equation (12) is

$$r=\beta+\theta g-0.5\gamma\sigma^2(\theta-1). \qquad (13)$$

Keeping β fixed at 0.05, as Obstfeld (1994) raised θ from 2 to 20, he implicitly attributed a value for the average after-tax return to capital between 9% and 49%

15. If $\gamma=\theta$, then equation (12) becomes equation (6).

per annum, well beyond the values in the data. As a result, his calculations heavily discounted the future costs arising from persistent shocks, which explains why he found that going from the Lucas to the Hall consumption models had little effect on the costs of fluctuations.

4. Statistical Models of Consumption

The key inputs into the costs of fluctuations are the volatility and the persistence of consumption. Both of these are notoriously difficult to estimate using the small U.S. post-war sample. This section pursues alternative approaches to do so.

4.1. Which Process for Consumption? Lucas versus Hall

At one extreme, consumption can have zero persistence (the Lucas model) or infinite persistence (the Hall model). Because these two models impose a rigid structure on the time-series of consumption, one can test which best describes the data.

Table 2 shows the results from different tests of the null hypothesis that consumption has a unit root: the original (augmented) test of Dickey and Fuller (1979), the alternative due to Phillips and Perron (1988), the point-optimal test of Elliott, Rothenberg, and Stock (1996), and finally the modified Phillips–Perron (MZ_t), point-optimal (MP_t), and Barghava statistic (MSB) tests combined with a modified Schwarz criteria to select the lag length. These last three tests were suggested by Ng and Perron (2001) in order to account for size distortions if the underlying data process is stationary. The results are clear: The null hypothesis corresponding to the Hall process is never rejected at the 5% significance level.

TABLE 2. Statistical tests of whether consumption has a unit root.

Test	Test statistic	5% critical value	Decision
Null hypothesis: Unit Root			
Dickey–Fuller	−1.88	−3.49	not rejected
Phillips–Perron	−1.74	−3.49	not rejected
Elliott–Rothenberg–Stock	10.95	5.71	not rejected
Ng–Perron:			
MZ_t	−2.02	−2.91	not rejected
MSB	0.24	0.17	not rejected
MP_t	10.83	5.48	not rejected
Null hypothesis: Stationarity			
Kwiatkowski et al.	0.17	0.15	rejected

Notes: The modified Schwarz criterion of Ng and Perron (2001) with a maximum lag of 10 selected the lag length of the regressions. For the Phillips–Perron and the Kwiatkowski et al. tests, I estimated the spectral density at frequency zero with a Bartlett kernel.

The last row of the table presents the result of a test by Kwiatkowski et al. (1992), of the null hypothesis that consumption is trend stationary. The data reject this hypothesis at the 5% significance level.

A simple way to understand why the data clearly favor the Hall process over the Lucas process is to nest both models in a single regression equation:

$$c_t - c_{t-1} = \text{constant} + u_t - \xi u_{t-1}, \tag{14}$$

where u_t is a residual. The Lucas process imposes the restriction $\xi = 1$, whereas the Hall process requires that $\xi = 0$. The 1947–2003 U.S. data produces an estimate of ξ of -0.36 with a standard error of 0.13. Not only is the estimate lower than one, it is not even positive—thus the strong statistical rejection of the Lucas model. However, note that although the Hall model is closer to the data, it is also rejected at the 5% significance level. Consumption growth is positively serially correlated, a fact that has inspired most modern research on consumption.[16] Fitting the facts requires richer models of consumption dynamics; the rest of this section investigates different possibilities.

4.2. Estimating the Persistence of U.S. Consumption

A statistical model for consumption that is more general than either the Lucas or the Hall models is the AR(1) in equation (3), where η need not be either one or zero. This way, persistence does not have to be zero or infinity, but it can be anywhere in between.

A naive application of this model would be to estimate η by least squares and, if the estimate is below 1, apply the formula in equation (4), as I did in Panel (B) of Table 1. However, it is well understood that for very persistent series like consumption, the least-squares estimate of η is downward-biased. For example, if the true model is a random walk, then the least-squares estimate of η will be below 1 with a probability of 68%. Given how steeply costs increase with η when it is close to 1, this can lead to severely underestimating the costs of fluctuations.

The most popular way to deal with this problem is to model η as lying within a circle of radius c/n around 1, where n is the size of the available sample. The estimate of the new parameter c (a "Pitman" drift) has a distribution that can be characterized using local-to-unity asymptotics (Stock 1994). Because deterministic formulae link c to η and in turn to λ, this characterizes the distribution of the estimate of the costs of fluctuations.

In the data, the confidence intervals for η include a large region above one. The formula in equation (2) would then imply that the costs of fluctuations are

16. See Fuhrer (2000) and Reis (2006) for two alternative models that try to account for this positive serial correlation, either by appealing to habits or to costs of processing information.

estimated to be infinity with a probability of more than 30%. This result arises because forecast error variances far ahead shoot quickly to infinity. This highlights one weakness of directly applying the formula in equation (2) if consumption follows an explosive process. The estimate of the costs of fluctuations in this case is dominated by estimates of the variability of consumption at horizons very far ahead, well above the size of the finite sample in which they were estimated.

The local-to-unity model suggests a natural way to deal with this issue. That model assumes that as the sample size increases, consumption becomes closer to a random walk; likewise, one can calculate the costs of fluctuations assuming that after the sample horizon, the forecast error variance is indistinguishable from that of a random walk. Focusing for now on the case of log utility, one estimator that formalizes this suggestion is

$$\hat{L} = 0.5(1 - e^{g-r}) \left[\sum_{t=0}^{n} e^{(g-r)t} \hat{v}(c_t) + \sum_{t=n+1}^{\infty} e^{(g-r)t} [\hat{v}(c_n) + \hat{v}(c_1)(t-n)] \right], \tag{15}$$

where $\hat{v}(c_t)$ is the least squares estimator of the the forecast error variance t steps ahead. This estimator replaces $\hat{v}(c_t)$ for horizons that exceed the size of the sample, by the n^{th} step-ahead forecast error variance for a random walk.[17] As $n \to \infty$, this estimator coincides with the exact value of the costs of fluctuations: $\hat{L} \to \ln(1 + \lambda)$. In a finite sample, under the maintained local-to-unity model, this is the estimator that is within $1/n$ of the costs of fluctuations.[18]

In the AR(1) model, straightforward but tedious algebra shows that, using the approximation $1 + c/n = \exp(c/n) + O(1/n^2)$,

$$\hat{L} = 0.5\hat{\sigma}^2 e^{g-r} \left[\sum_{t=0}^{n} e^{(g-r+2c/n)t} + \frac{e^{(g-r)n}}{1 - e^{g-r}} \right], \tag{16}$$

where $\hat{\sigma}$ is the least-squares estimate of the standard error of shocks. It is easy to show that as $n \to \infty$, $\hat{\sigma}^2 \to \sigma^2$. Applying the functional central limit theorem,

$$\frac{1}{n} \sum_{t=0}^{n} e^{(g-r-2c/n)t} \Rightarrow \int_0^1 e^{(g-r+2U)s} ds, \tag{17}$$

where $\Rightarrow$ denotes weak convergence. The random variable U is

$$U = \frac{\int_0^1 J(s) dW(s)}{\int_0^1 J(s)^2 ds},$$

17. For a random walk, $\mathrm{Var}(c_t) = \sigma^2 t$, so $\mathrm{Var}(c_t) = \mathrm{Var}(c_n) + \mathrm{Var}(c_1)(t - n)$ for $t > n$.

18. This approach has a close relative in Phillips's (1998) construction of confidence intervals for far-ahead impulse responses in the local-to-unity model.

where $J(\cdot)$ is an Orstein–Uhlenbeck process $dJ(s) = cJ(s)ds + dW(s)$ and $W(\cdot)$ is a standard Brownian motion. The continuous mapping theorem then implies that

$$\frac{\hat{L}}{n} \Rightarrow 0.5\sigma^2 e^{g-r} \frac{e^{g-r+2U} - 1}{g - r + 2U}, \tag{18}$$

which fully describes the asymptotic distribution of the estimate of the costs of fluctuations.

According to this asymptotic result, the least-squares estimate of the costs of fluctuations is not only an inconsistent estimate of the true costs, but moreover, it converges to a random variable. The reason is that as n grows, the least-squares estimation errors persist for longer rather than dying off. This implies that the estimates in Panel (B) of Table 1 were downward-biased. Yet using the formula in equation (18), constructing median-unbiased estimates and confidence intervals for the costs of fluctuations is possible.[19] Appendix A.5 extends the calculations in this section in two directions. First, it considers the case when relative risk aversion is different from one. Second, it extends the statistical model to the Dickey–Fuller regression form:

$$\Delta c_t = \kappa_0 + \kappa_1 t + \varrho c_{t-1} + \sum_{j=1}^{k} \psi_j \Delta c_{t-j} + u_t. \tag{19}$$

Now, it is the largest autoregressive root that is modeled as $1 + c/n$. This allows for a more flexible characterization of log consumption, as a $(k + 1)$-order autoregressive process with a drift and a time trend.

Table 3 presents median-unbiased estimates and 90% confidence intervals for the estimated costs of fluctuations if consumption dynamics are described by equation (19). The costs are now much higher than the naive estimates in Table 1. They range from 0.2% to 3.2% of per capita consumption and even the lower bounds of the confidence intervals are higher than those in the Panel (B) of Table 1. According to these calculations, society substantially dislikes the current variability in consumption.

4.3. Parametric Unrestricted Estimates

Because the evidence in Section 4.1. suggested that consumption is not stationary, I now use this as a starting point to build alternative statistical models of consumption and its persistence. The data do not reject the null hypothesis that the

19. The distribution of U is non-normal and depends on the (unknown) value of c. It therefore requires many numerical simulations to characterize this distribution for each value of c. Stock (1991) has already done the work of tabulating the distribution of U. Because $\hat{L}/n$ increases monotonically with U, one can use his tables to construct confidence intervals for the estimates of the costs of fluctuations.

TABLE 3. The costs of fluctuations when consumption is persistent.

Panel A: Costs in percentages of annual per capita consumption			
	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.21%	0.63%	1.05%
	(0.19–0.21)	(0.57–0.64)	(0.95–1.07)
$\rho = 0.02$	0.31%	0.95%	1.58%
	(0.29–0.32)	(0.87–0.96)	(1.46–1.61)
$\rho = 0.01$	0.63%	1.90%	3.19%
	(0.60–0.64)	(1.81–1.92)	(3.03–3.22)
Panel B: Costs in annual per capita 2003 dollars			
	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	\$49	\$147	\$246
	(44–50)	(133–150)	(223–256)
$\rho = 0.02$	\$74	\$222	\$371
	(68–75)	(204–225)	(341–377)
$\rho = 0.01$	\$148	\$446	\$747
	(140–149)	(423–450)	(709–755)

Notes: Each cell shows the median unbiased estimate and, in parentheses, the 90% confidence interval. The Ng and Perron (2001)–modified BIC picked the autoregression's order.

first difference of consumption is stationary. (The unit root tests are not reported here for brevity.) Wold's theorem states that a stationary series has a moving average representation, so if consumption is integrated of order one, then a general statistical model for consumption is

$$\Delta c_t = \text{constant} + A(L)u_t, \tag{20}$$

where $\Delta c_t = (1 - L)c_t$ and $A(L) = \sum_{i=0}^{\infty} a_i L^i$, and L is the lag operator $L^i u_t = u_{t-i}$.

If consumption follows this process, the costs of fluctuations in equation (2) become

$$\ln(1+\lambda) = \begin{cases} 0.5\sigma^2(1-e^{g-r})\left(\sum_{t=1}^{\infty} e^{(g-r)t} \sum_{j=0}^{t-1}\sum_{i=0}^{j} a_i^2\right) & \text{if } \gamma = 1, \\ (\gamma-1)^{-1}\ln\Big[(1-e^{g-r}) & \\ \quad \times \left(1+\sum_{t=1}^{\infty} e^{(g-r)t} e^{0.5\sigma^2\gamma(\gamma-1)\sum_{j=0}^{t-1}\sum_{i=0}^{j} a_i^2}\right)\Big] & \text{if } \gamma \neq 1. \end{cases} \tag{21}$$

It is impossible to estimate the infinite number of parameters a_i with a finite number of observations. However, it has long been known that an ARMA model

$$B(L)\Delta c_t = \text{constant} + C(L)\varepsilon_t, \tag{22}$$

where $B(L)$ and $C(L)$ are lag polynomials of low order, typically provides a good approximation to the dynamics of most macroeconomic series. Given estimates of the ARMA model, one can easily recover the parameters a_i using the relation $A(L) = B(L)^{-1}C(L)$.

Estimating equation (22) requires choosing the order of $B(L)$ and $C(L)$. I restricted the range of admissible models to a maximum of three AR and/or MA parameters. ARMA processes with many parameters are notoriously difficult to estimate and the experience with ARMA modelling has been that low-order ARMA processes typically have a superior forecasting performance. I estimated the 16 admissible models by maximum likelihood.[20] To pick between them, I used the Bayesian information criterion (BIC). This criterion picks the model with the highest likelihood, while imposing a penalty that increases with the number of parameters being estimated. One advantage of the BIC is that, as the sample size goes to infinity, it consistently picks the true underlying model. The BIC picked the ARMA(2,2) as the best model, followed by the ARMA(1,0) and by the ARMA(0,1).

Table 4 shows the costs of eliminating fluctuations in consumption for these three statistical models. The first conclusion to take from the table is that the estimates are all larger than the corresponding estimates in Panel (C) of Table 1. The positive serial correlation in consumption growth implies that shocks propagate by more over time than what the Hall model predicted. A second conclusion is that across the three empirical consumption processes, the estimates of the costs of fluctuations are roughly similar. The results are robust in the sense that moving between models that fit the data almost equally well does not drastically affect the estimates. This leads to the third conclusion: The costs of fluctuations are approximately between 0.5% and 5% of per capita consumption, similar to the estimates in Table 3.

4.4. Non-Parametric Unrestricted Estimates

Although, so far, I have been focusing on the persistence of consumption, the key empirical inputs into the formula for the costs of fluctuations in equation (2) are the forecast error variances of consumption. Until now, I have estimated these by fitting parametric models to the observations of consumption. A natural alternative is to estimate the forecast error variances directly imposing as little structure as possible on the model of consumption.

Because these variances are conditional on information at time zero, then $\mathrm{Var}(c_t) = \mathrm{Var}(c_t - c_0)$. It is difficult to estimate these without specifying what

20. One important concern with estimating ARMA models is that the likelihood functions are often multi-peaked or nearly flat for a wide range of parameter values, so numerical procedures can converge on incorrect estimates. To safeguard against this possibility, I plotted the likelihood functions, examined their gradients at the proposed optima, and started the numerical maximizations from different initial values.

TABLE 4. The costs of fluctuations estimated by ARMA models.

	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
Panel A: Estimated ARMA (2,2) model $(1 - 0.66L - 0.32L^2)\Delta c_t = (1 + 1.03L + 0.56L^2)u_t, \sigma_u = 0.011$			
$\rho = 0.03$	0.31%	0.94%	1.60%
	(\$72)	(\$219)	(\$375)
$\rho = 0.02$	0.47%	1.43%	2.47%
	(\$109)	(\$334)	(\$579)
$\rho = 0.01$	0.94%	2.93%	5.33%
	(\$220)	(\$687)	(\$1248)
Panel B: Estimated ARMA (1,0) model $(1 - 0.34L)\Delta c_t = u_t, \sigma_u = 0.010$			
$\rho = 0.03$	0.40%	1.23%	2.13%
	(\$94)	(\$288)	(\$498)
$\rho = 0.02$	0.61%	1.89%	3.33%
	(\$144)	(\$442)	(\$780)
$\rho = 0.01$	1.25%	3.94%	7.40%
	(\$292)	(\$923)	(\$1734)
Panel C: Estimated ARMA (0,1) model $\Delta c_t = (1 + 0.36L)u_t, \sigma_u = 0.011$			
$\rho = 0.03$	0.34%	1.02%	1.76%
	(\$79)	(\$240)	(\$412)
$\rho = 0.02$	0.51%	1.56%	2.73%
	(\$120)	(\$366)	(\$638)
$\rho = 0.01$	1.03%	3.23%	5.92%
	(\$242)	(\$752)	(\$1388)

Notes: Each cell shows the per capita costs of fluctuations as a fraction of consumption and, in parentheses, in 2003 U.S. dollars.

the conditioning information at time 0 is. However, doing so is close to specifying a parametric model for consumption, precisely what this section is trying to avoid. I overcome this dilemma by estimating the unconditional variance of the tth difference in log consumption. The conditional and unconditional variances will be the same in the case of the AR(1); otherwise, the unconditional variances will be higher. The estimates in this section therefore provide non-parametric upper bounds on the costs of fluctuations.

Cochrane (1987) showed that the unconditional variance of the tth difference in consumption equals

$$t\sigma_{\Delta c}^2 \left(1 + 2\sum_{j=1}^{t-1} \frac{t-j}{t} R_j\right). \tag{23}$$

R_j is the jth-order autocorrelation of the first difference of consumption; $\sigma_{\Delta c}^2$ is its variance. In parentheses is the Bartlett estimator of the spectrum of the first-difference of consumption at frequency zero using a lag window of length t. The sample autocorrelations and variance of the first difference of consumption provide consistent estimates of these moments.[21]

One difficulty is that it is impossible to compute the variance of the tth difference in consumption if t is larger than the sample size. Even if t is smaller than n, as long as it is close to it, the estimator of R_t will be using only a few observations. I tackle this problem in the same way that I did earlier when deriving the asymptotic distribution of the costs of fluctuations. I use an estimator like $\hat{L}$ in equation (15), with the only difference that the first sum now includes terms only up to a fraction of n. This way, the estimator only requires computing the variances of consumption differences up to a fraction of the sample. As before, this estimator asymptotically converges to the true costs of fluctuations and it provides a good approximation in a finite sample if consumption is very persistent.

Table 5 contains the new estimates of the costs of fluctuations. From Panels (A) to (C), I use increasing fractions of the sample, from 25% to 50% to 75%. The costs of fluctuations from using this approach are typically in between the random walk estimates and the larger estimates using ARMA models. They are all larger than the Lucas benchmark of 0.05%.

5. Economic Models of Consumption

I now turn to economic models of consumption, that is, models in which society optimally uses savings to choose a path for consumption over time.

5.1. The Components of the Models

I focus on a pervasive model of consumption and fluctuations, the neoclassical stochastic growth model. Society's welfare is given by the value function of a representative consumer $V_\tau(K, A)$ defined by

$$V_\tau(K, A) = \max_{\{C_t\}} \quad E\left[\sum_{t=0}^{\infty} e^{-\beta t} u(C_t)\right]$$

$$\text{subject to:} \quad K_{t+1} = A_t^{1-\alpha} K_t^\alpha + (1-\delta)K_t - C_t.$$

Here, A_t is stochastic productivity, the log of which follows the process

$$a_t = g\varphi + g(1-\varphi)t + \varphi a_{t-1} + e_t, \quad \text{with } e_t \sim N(0, \tau^2), \tag{24}$$

21. I multiply the estimator by $n/(n-t+1)$ to improve its performance in a small sample (Cochrane 1987).

TABLE 5. The costs of fluctuations estimated using the unconditional variances.

Panel A: Estimating correlations of order up to 25% of the sample			
	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.25%	0.76%	1.29%
	(\$59)	(\$179)	(\$303)
$\rho = 0.02$	0.36%	1.10%	1.87%
	(\$85)	(\$257)	(\$440)
$\rho = 0.01$	0.68%	2.10%	3.68%
	(\$160)	(\$492)	(\$863)
Panel B: Estimating correlations of order up to 50% of the sample			
	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.26%	0.78%	1.33%
	(\$61)	(\$183)	(\$311)
$\rho = 0.02$	0.37%	1.12%	1.91%
	(\$86)	(\$262)	(\$447)
$\rho = 0.01$	0.69%	2.12%	3.71%
	(\$161)	(\$496)	(\$871)
Panel C: Estimating correlations of order up to 75% of the sample			
	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.16%	0.48%	0.81%
	(\$37)	(\$113)	(\$189)
$\rho = 0.02$	0.22%	0.68%	1.15%
	(\$52)	(\$159)	(\$271)
$\rho = 0.01$	0.48%	1.47%	2.60%
	(\$112)	(\$346)	(\$610)

Notes: Each cell shows the per capita costs of fluctuations as a fraction of consumption and, in parentheses, in 2003 U.S. dollars.

which is consistent with data on Solow residuals. As before, g measures the expected growth rate of consumption in the model, and φ measures the persistence of productivity shocks. The two state variables are K_t and A_t and the economy starts from some K_0 and A_0. I index the value functions by the standard deviation of the shocks τ.

The depreciation rate δ is set at 0.05, although close alternatives have little effect on the costs of fluctuations. The parameter α is set at 0.36 to match the U.S. share of payments to capital. Section 2 already inspected the $\alpha = 1$ case, and a previous draft of this paper considered setting α at 0.75, consistent with a broad view of capital that includes human capital and with the typical estimates of conditional convergence. Because a higher α implies that diminishing returns only slowly set in, it leads to more persistent movements in consumption, and raises the costs of fluctuations. Going from 0.36 to 0.75 typically doubled the costs of fluctuations.

Different models correspond to different choices of γ, $\rho = \beta + (\gamma - 1)g$, and φ. The focus is on this last parameter because it is the key determinant of the persistence of consumption. I consider two cases, both consistent with the data: $\varphi = 0.9$ so productivity and consumption are trend-stationary, and $\varphi = 1$ so productivity and consumption are non-stationary.

Aside from persistence, the other key determinant of the costs of fluctuations is volatility, in this model determined by τ. Solow residuals are not very helpful at pinning down this parameter because of the noise in measuring productivity. Consistent with the general approach in this paper, I calibrate this parameter to match the properties of consumption. Namely, I pick τ so that the model matches either the standard deviation of log consumption (for the stationary model) or its first difference (for the non-stationary model).

5.2. *Solving for the Costs of Fluctuations*

In the counterfactual economy without fluctuations, productivity is equal to $E[A_t]$ and the value function is denoted by $W(K, A)$.[22] For $\gamma = 1$, the definition of the costs of fluctuations in equation (1) implies that

$$\ln(1+\lambda) = (1 - e^{-\rho})(W(K, A) - V_\tau(K, A)).$$

The case $\gamma \neq 1$ is similar and left to Appendix A.6.

When productivity is stationary, the value functions can be re-expressed in terms of the stochastic components $\hat{C}_t = C_t e^{-gt}$, $\hat{K}_t = K_t e^{-gt}$ and $\hat{a}_t = a_t - gt$ through the relation

$$V_\tau(K, A) = v_\tau(\hat{K}, \hat{a}) + \frac{g}{(1 - e^{-\rho})(e^{\rho} - 1)},$$
$$W(K, A) = w(\hat{K}, \hat{a}) + \frac{g}{(1 - e^{-\rho})(e^{\rho} - 1)},$$

where a little work shows that

$$v_\tau(\hat{K}, \hat{a}) = \max_{\hat{C}} \quad \left[\ln(\hat{C}) + e^{-\rho} E[v_\tau(\hat{K}', \varphi\hat{a} + e')]\right],$$
$$\text{subject to:} \quad \hat{K}' = e^{-g}(e^{(1-\alpha)\hat{a}} \hat{K}^{\alpha} + (1-\delta)\hat{K} - \hat{C}_t),$$

and

$$w(\hat{K}, \hat{a}) = \max_{\hat{C}} \quad \left[\ln(\hat{C}) + e^{-\rho} w(\hat{K}', \varphi\hat{a})\right],$$
$$\text{subject to:} \quad \hat{K}' = e^{-g}\left(e^{(1-\alpha)(\hat{a}+0.5\tau^2/(1-\varphi^2))} \hat{K}_t^{\alpha} + (1-\delta)\hat{K}_t - \hat{C}_t\right).$$

22. The expectation is conditional at date 0, so to eliminate the effect of initial conditions, I assume that the date t at which welfare is evaluated is far into the future.

From the statement of the value functions, one can see that

$$w(\hat{K}, \hat{a}) = \frac{0.5\tau^2}{(1-\varphi^2)(1-e^{-\rho})} + v_0(\hat{K}, \hat{a}).$$

The costs of fluctuations are then

$$\ln(1+\lambda) = \frac{0.5\tau^2}{1-\varphi^2} + (1-e^{-\rho})(v_0(\hat{K}, \hat{a}) - v_\tau(\hat{K}, \hat{a})). \tag{25}$$

The calculation of the costs of fluctuations then boils down to solving for the value function $v_\tau(\hat{K}, \hat{a})$. I do this using the algorithm in Fackler (2005) that approximates the expectations by Gaussian quadrature, approximates the unknown functions by Chebyshev polynomials, and solves the system by collocation and Broyden's method. I solve these dynamic programs for several τ including 0 and simulate paths for consumption for each, picking the one that produces a standard deviation of consumption that matches the U.S. data.

With non-stationary productivity, I am only able to solve the model when $\gamma = 1$. These models do not have a steady state, making numerical implementation hard, with the exception of the $\gamma = 1$ case, for which Christiano (1988) found a stationary-inducing transformation. Letting $\tilde{K}_t = K_t/A_{t-1}$ and $\tilde{C}_t = C_t/A_t$, similar algebra shows that

$$V_\tau(K, A) = v_\tau(\tilde{K}) + \frac{a}{1-e^{-\rho}},$$
$$W(K, A) = v_0(\tilde{K}) + \frac{a + 0.5\tau^2/(e^\rho - 1)}{1-e^{-\rho}},$$

where now

$$v_\tau(\hat{K}) = \max_{\hat{C}} \{\ln(\hat{C}) + e^{-\rho} E[v_\tau(e^{-\alpha(g+e)}\hat{K}^\alpha + (1-\delta)e^{-(g+e)}\hat{K} - \hat{C}_t)]\}.$$

From this, it follows that the costs of fluctuations are

$$\ln(1+\lambda) = \frac{0.5\tau^2}{e^\rho - 1} + (1-e^{-\rho})(v_0(\tilde{K}) - v_\tau(\tilde{K})). \tag{26}$$

The value functions for this case are solved just as in the stationary case.

Finally, as the expressions in equations (25) and (26) make clear, the estimates of the costs of fluctuations depend on the $(\hat{K}, \hat{a})$ point at which they are evaluated. This point matters, for instance because the curvature of the value function may change and so will society's aversion to fluctuations. I consider three alternatives. The first two evaluate the costs of fluctuations at $\hat{a} = 0$ and at a capital stock equal to either its non-stochastic steady state, or its average in the stochastic model. The third measure instead integrates over the stationary distributions of $\hat{a}$ and the capital stock predicted by the model, calculating an unconditional expectation of the costs of fluctuations.

5.3. *Estimates of the Costs of Fluctuations in the Economic Model*

Table 6 presents the estimated costs of fluctuations in the stationary model and Table 7 presents those in the non-stationary model. As expected, raising γ typically raises the costs of fluctuations. Changes in ρ typically have modest effects on the costs of fluctuations in the stationary case, in contrast with what happened in the statistical models. The reason seems to be that the agent reacts strongly to different discount rates by changing her optimal allocation of consumption over time. In the simulations, the average capital stock was quite different across different choices of ρ. In the non-stationary case, the shocks are permanent so there is little that intertemporal smoothing can do about them and lowering ρ significantly raises the costs of fluctuations.

Evaluating the costs of fluctuations at different points gives slightly different results. The difference between the costs at the mean capital stock versus the steady-state level are typically negligible, never larger than 0.02%. The unconditionally expected costs of fluctuations can be as much as 0.04% larger than the costs at any of the other points, because the distribution includes some points at which the curvature of the value function can be quite higher.

The focus of this paper is on the contrast between Tables 6 and 7. If productivity is stationary, the costs of fluctuations are small, typically between 0.05% and 0.10% of consumption. In comparison, the estimates in Table 7 are much larger, between 0.1% and 0.3% of consumption. Economic models therefore yield the same lesson as statistical models: Persistence is a key determinant of the costs of fluctuations. Unlike the statistical models though, the costs of fluctuations are never very large. The highest estimate in the table is that each person in the United States would be willing to pay $77 to eliminate fluctuations in consumption. This is likely explained by the representative consumer accumulating a buffer stock of savings to self-insure against some of the income fluctuations.

The class of models analyzed in this section is just one among many different possibilities. Aside from generating estimates of the costs of fluctuations that are interesting in their own right, these models served a dual purpose. First, they showed how to calculate the costs of fluctuations within economic models that take the precautionary savings and investment risk effects into account. Second, they showed that the model's predicted persistence of consumption is a key determinant of the costs of fluctuations. This opens the door to estimating the costs of fluctuations in richer models that have other features to fit the business cycle facts (e.g., as in Otrok 2001) and other mechanisms to propagate shocks over time aside from investment, such as nominal rigidities or credit market frictions.

6. Interpretation of the Results and Conclusion

This paper re-examined the estimation of the costs of fluctuations. It showed that the properties of the stochastic process describing consumption, and especially the

TABLE 6. The costs of fluctuations with stationary productivity.

Panel A: Evaluated at the non-stochastic steady state capital stock	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.03%	0.05%	0.09%
	($6)	($13)	($22)
$\rho = 0.02$	0.04%	0.06%	0.10%
	($10)	($13)	($23)
$\rho = 0.01$	0.03%	0.08%	0.10%
	($8)	($19)	($24)
Panel B: Evaluated at the mean capital stock with fluctuations	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.05%	0.06%	0.10%
	($12)	($13)	($24)
$\rho = 0.02$	0.03%	0.06%	0.10%
	($8)	($15)	($23)
$\rho = 0.01$	0.02%	0.08%	0.10%
	($4)	($19)	($24)
Panel C: Unconditional expectation over capital and productivity	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
$\rho = 0.03$	0.04%	0.10%	0.14%
	($10)	($23)	($32)
$\rho = 0.02$	0.03%	0.06%	0.10%
	($8)	($15)	($23)
$\rho = 0.01$	0.02%	0.08%	0.10%
	($4)	($19)	($24)

Notes: Each cell shows the per capita costs of fluctuations as a fraction of consumption and, in parentheses, in 2003 U.S. dollars per capita.

persistence of shocks, are a key determinant of the costs of fluctuations. Although the assumptions made by Lucas (1987) are decisively rejected by the data, I have found that if consumption is only mildly persistent, the estimated costs of fluctuations are close to those that Lucas estimated. However, the evidence suggests that consumption fluctuations are more persistent than this, and as persistence increases, the costs of fluctuations rise substantially. For instance, if consumption

TABLE 7. The costs of fluctuations with non-stationary productivity.

	At the non-stochastic steady-state capital stock	At the mean capital stock	Unconditional expectation
$\rho = 0.03$	0.10%	0.10%	0.10%
	($24)	($23)	($24)
$\rho = 0.02$	0.16%	0.15%	0.17%
	($38)	($35)	($40)
$\rho = 0.01$	0.32%	0.32%	0.32%
	($76)	($75)	($76)

Notes: Each cell shows the per capita costs of fluctuations as a fraction of consumption and, in parentheses, in 2003 U.S. dollars per capita.

is a random walk, as some theories suggest and the data does not reject, the costs of fluctuations are fifty to one hundred times larger than what Lucas estimated. The statistical models that best fit the data and the economic models that account for the effect of fluctuations on precautionary savings lead to large estimates of the costs of fluctuations, typically between one and two orders of magnitude larger than Lucas's benchmark.

As discussed in the Introduction, this paper does not take a stand as to whether the fluctuations behind the estimates in this paper correspond to business cycles. If business cycles are transitory short-lived deviations of consumption away from a stable trend, as defined by for instance the use of band-pass filters, this paper suggests that the costs of business cycles are small. However, there is an alternative view of business cycles that dates back at least to Burns and Mitchell (1946) and which defines cycles as a set of regularities in the co-movement of macroeconomic series. Campbell and Mankiw (1987) found that output fluctuations in the United States are actually very long-lived and Kydland and Prescott (1982) showed that a calibrated real business cycle model driven by non-stationary productivity shocks and consumption fits the U.S. data well. Under this view of business cycles, the welfare costs may be quite large.

Likewise, this paper does not take a stand on whether the fluctuations are efficient or whether policy-makers can eliminate them. Still, costs between 0.5% and 5% of consumption are significant. To put them into perspective, in 2003 the total amount spent by the U.S. federal government in unemployment and medical insurance was $53 billion, or 0.8% of consumption; the amount spent in consumption by the federal government excluding national defense was $223 billion (3.3%); the amount spent in health coverage for low-income families through the Medicaid program was $265 billion (3.9%).[23] Even if policy could eliminate only a part of these costs, the benefits would be quite significant. They are certainly smaller than raising the economy's growth rate by 1% but this comparison is only fair if it is as easy to raise the growth rate as it is to dampen fluctuations. There is little evidence of economists' success at affecting growth (Easterly 2002) but some shreds of evidence that advances in economic knowledge may have led to policies that stabilized the economy (Romer and Romer 2002). Instead, dampening even only part of consumption fluctuations leads to gains in the same range as other potentially feasible policies that Lucas (2003) discussed, like lowering inflation from 10% to zero (gain of 1%) or eliminating capital income taxes (2%–4% gain).

Whichever view one takes of fluctuations, the calculations in this paper have at least provided the tools to estimate their costs under different scenarios. They put upper bounds on how much dampening fluctuations can improve welfare,

23. Source: National Income and Product Accounts, Tables 3.9.5 and 3.12.

bounds that are large enough to motivate further work on figuring out how much of these are inefficient and how much policy can achieve.[24]

Appendix

A.1. The Costs of Fluctuations in Statistical Models

For the case $\gamma = 1$, the definition of the costs of fluctuations in equation (1) and of the counterfactual in statistical models imply that

$$\ln(1+\lambda) + (1-e^{-\beta}) \sum_{t=0}^{\infty} e^{-\beta t} E[c_t] = (1-e^{-\beta}) \sum_{t=0}^{\infty} e^{-\beta t} (E[c_t] + 0.5\,\mathrm{Var}(c_t)). \tag{A.1}$$

This result used the log-normality of C_t to evaluate $\ln(E[C_t])$. Rearranging and substituting β for ρ gives the first expression in equation (2). For $\gamma \neq 1$, log-normality of consumption implies that

$$E\big[C_t^{1-\gamma}\big] = E\big[C_t\big]^{1-\gamma} e^{0.5\gamma(\gamma-1)\,\mathrm{Var}(c_t)}.$$

Similar rearrangements lead to the second expression in equation (2).

A.2. The Costs of Fluctuations in the AR(1) Model

For a stationary AR(1), $\mathrm{Var}(c_t) = \sigma^2(1-\eta^{2t})/(1-\eta^2)$ for $t \geq 1$. When $\gamma = 1$, evaluating the sum in equation (2) shows that

$$\ln(1+\lambda) = \frac{0.5\sigma^2}{e^{\rho} - \eta^2}. \tag{A.2}$$

Using the approximations $e^{\rho} - 1 \cong \rho$ and $\ln(1+\lambda) \cong \lambda$ gives the result.

For the case when $\gamma \neq 1$, approximate

$$\ln(1+\lambda) = \frac{1}{\gamma-1} \ln\left[(1-e^{\rho}) \sum_{t=0}^{\infty} e^{-\rho t} e^{0.5\gamma(\gamma-1)\sigma^2(1-\eta^{2t})/(1-\eta^2)}\right] \tag{A.3}$$

around $\sigma^2 = 0$ using a first-order Taylor expansion. Terms of order σ^4 or higher are tiny in the data, so this involves little error. This leads immediately to the same expression as in the log case, but now multiplied by γ:

$$\ln(1+\lambda) \cong 0.5\gamma\sigma^2/(e^{\rho} - \eta^2).$$

24. Galí, Gertler, and Lopez-Salido (2007) recently proposed a promising method to disentangle efficient from inefficient fluctuations and measure their cost.

Similar approximations to before give the final result.

A.3. The Costs of Fluctuations in the Benchmark Economic Model

The Euler equation for the problem in equations (6)–(7) is

$$C_t^{-\gamma} = e^{-\beta} E_t\left[R_{t+1} C_{t+1}^{-\gamma}\right]. \tag{A.4}$$

Guess that consumption is linear in wealth, $C_t = \pi R_t K_t$, with a coefficient π to be determined. The budget constraint implies that

$$\frac{C_{t+1}}{C_t} = \frac{R_{t+1} K_{t+1}}{R_t K_t} = R_{t+1}(1 - \pi). \tag{A.5}$$

Using this result to replace for C_{t+1}/C_t in the Euler equation and the fact that R_{t+1} is log-normally distributed, equation (A.5) becomes

$$\gamma \ln(1 - \pi) = (1 - \gamma)r - \beta + 0.5\gamma(\gamma - 1)\sigma^2. \tag{A.6}$$

This expression does not depend on any state variable, which confirms the initial guess. Using the definition of ρ in equations (A.6) and (A.5) to substitute π out, gives the expressions in equations (8)–(10).

The costs of fluctuations for $\gamma \neq 1$ solve the equation

$$(1 + \lambda)^{1-\gamma}(1 - e^{g-r})^{1-\gamma} \sum_{t=0}^{\infty} e^{[-\beta+(1-\gamma)(g-0.5\gamma\sigma^2)]t} =$$

$$(1 - e^{g-r+0.5(1-\gamma)\sigma^2})^{1-\gamma} \sum_{t=0}^{\infty} e^{[-\rho+(1-\gamma)(g+0.5(1-\gamma)\sigma^2)]t} \tag{A.7}$$

Use the definition of g in equation (8) to replace for β and obtain

$$(1 + \lambda)^{1-\gamma}(1 - e^{g-r})^{1-\gamma} \sum_{t=0}^{\infty} e^{(g-r)t} =$$

$$(1 - e^{g-r+0.5(1-\gamma)\sigma^2})^{1-\gamma} \sum_{t=0}^{\infty} e^{[g-r+0.5(1-\gamma)\sigma^2]t}. \tag{A.8}$$

Evaluating the sums and taking logs shows that

$$\ln(1 + \lambda) = \frac{\gamma}{\gamma - 1} \ln\left(\frac{e^{r-g} - e^{-0.5(\gamma-1)\sigma^2}}{e^{r-g} - 1}\right). \tag{A.9}$$

A first-order Taylor approximation of this expression on σ^2, equation (9), and the approximations $e^{\rho} - 1 \cong \rho$ and $\ln(1+\lambda) \cong \lambda$ gives the result in equation (11). The case when $\gamma = 1$ follows along the same steps.

A.4. The Costs of Fluctuations in the Epstein–Zin–Weil Model

Weil (1990) shows that optimal consumption is in equation (8) but now with

$$\theta g = r - \beta + 0.5(\theta - 1)\gamma\sigma^2. \tag{A.10}$$

The expected discounted utility with optimal consumption choices equals (up to a constant)

$$\frac{(R_0 K_0)^{1-\gamma}(1-e^{-\beta})^{(1-\gamma)/(1-\theta)}(1-e^{g-r})^{-\theta(1-\gamma)/(1-\theta)}}{(1-e^{-\beta})(1-\gamma)}. \tag{A.11}$$

With the preferences in equation (12), without fluctuations, discounted utility equals (up to the same constant)

$$\frac{(1-e^{-\beta})^{(1-\gamma)/(1-\theta)}(R_0 K_0)^{1-\gamma}(1-e^{g-r-(\theta-1)0.5\gamma\sigma^2/\theta})^{-\theta(1-\gamma)/(1-\theta)}}{(1-e^{-\beta})(1-\gamma)}. \tag{A.12}$$

Given the definition of the costs of fluctuations, $(1+\lambda)^{1-\gamma}$ equals the ratio of the first terms in equations (A.12) and (A.11). After cancelling some terms and taking logs, this equals

$$\ln(1+\lambda) = \frac{\theta}{\theta - 1}\ln\left(\frac{e^{r-g} - e^{-0.5(\theta-1)\gamma\sigma^2/\theta}}{e^{r-g} - 1}\right). \tag{A.13}$$

Finally, note that a linear approximation of the right-hand side of equation (A.13) in σ^2 around zero is equal to a linear approximation of the right-hand side of equation (A.9).

A.5. Asymptotic Distributions for the Extended Auto-Regressive Model

The first extension is to when $\gamma \neq 1$. The simplest way to do this is to approximate the definition of the costs of fluctuations in equation (2) around the point $\sigma^2 = 0$. This shows that up to terms that are $O(\sigma^4)$ the costs of fluctuations with $\gamma \neq 1$ just equal γ times the costs for the log utility case. In the data, the estimates of σ are typically tiny so the σ^4 terms being ignored are quantitatively insignificant.

The second extension is to the Dickey–Fuller regression. One change is that now $\varrho = 1+(c/n)(1-\sum_{j=1}^{k}\psi_j)$. Another change is that the distribution of ϱ is affected by the presence of the constant and the trend. Stock (1991) showed that

$$n(\hat{\varrho}-1) \Rightarrow \left(1-\sum_{j=1}^{k}\psi_j\right)\left[\left(\int_0^1 J^{\tau}(s)^2 ds\right)^{-1}\left(\int_0^1 J^{\tau}(s)dW(s)\right)+c\right], \tag{A.14}$$

where $J^{\tau}(s) = J(s)-\int_0^1(2-6r)J(r)dr - s\int_0^1(12r-6)J(r)dr$. The distribution of the estimate of the costs of fluctuations is otherwise similar to before.

A.6. The Costs of Fluctuations with Stationary Productivity

The main text dealt with the case $\gamma = 1$. Following very similar steps, with $\gamma \neq 1$ the definition of the costs of fluctuations in equation (1) implies

$$\ln(1+\lambda) = \frac{\ln(W(K,A)) - \ln(V_\tau(K,A))}{1-\gamma},$$

The value functions in terms of stochastic components are $V_\tau(K,A) = v_\tau(\hat{K},\hat{a})$ and $W(K,A) = w(\hat{K},\hat{a})$ with

$$\begin{aligned} v_\tau(\hat{K},\hat{a}) = \max_{\hat{C}} \quad & \left[\hat{C}^{1-\gamma}/(1-\gamma) + e^{-\rho}E[v_\tau(\hat{K}',\varphi\hat{a}+e')]\right] \\ \text{subject to:} \quad & \hat{K}' = e^{-g}(e^{(1-\alpha)\hat{a}}\hat{K}^{\alpha} + (1-\delta)\hat{K} - \hat{C}_t), \end{aligned}$$

and

$$\begin{aligned} w(\hat{K},\hat{a}) = \max_{\hat{C}} \quad & \left[\hat{C}^{1-\gamma}/(1-\gamma) + e^{-\rho}w(\hat{K}',\varphi\hat{a})\right] \\ \text{subject to:} \quad & \hat{K}' = e^{-g}\left(e^{(1-\alpha)(\hat{a}+0.5\tau^2/(1-\varphi^2))}\hat{K}_t^{\alpha} + (1-\delta)\hat{K}_t - \hat{C}_t\right) \end{aligned}$$

Now, we have that $w(\hat{K},\hat{a}) = \exp\{(1-\gamma)0.5\tau^2/(1-\varphi^2)\}v_0(\hat{K},\hat{a})$. The costs of fluctuations are then

$$\ln(1+\lambda) = \frac{0.5\tau^2}{1-\varphi^2} + \frac{\ln(v_0(\hat{K},\hat{a})) - \ln(v_\tau(\hat{K},\hat{a}))}{1-\gamma}.$$

References

Alvarez, Fernando, and Urban J. Jermann (2004). "Using Asset Prices to Measure the Cost of Business Cycles." *Journal of Political Economy*, 112, 1223–1256.

Andrews, Donald W. K. (1991). "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3), 817–858.

Angeletos, George-Marios (2007). "Uninsured Idiosyncratic Investment Risk and Aggregate Saving." *Review of Economic Dynamics*, 10(1), 1–30.

Atkeson, Andrew, and Christopher Phelan (1994). "Reconsidering the Costs of Business Cycles with Incomplete Markets." In *NBER Macroeconomics Annual 1994*, edited by Stanley Fisher and Julio J. Rotemberg. MIT Press.

Barlevy, Gadi (2004). "The Cost of Business Cycles under Endogenous Growth." *American Economic Review*, 94(4), 964–990.

Barro, Robert (2006). "On the Welfare Costs of Consumption Uncertainty." NBER Working Paper No. 12763.

Beaudry, Paul, and Carmen Pages (2001). "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance." *European Economic Review*, 45, 1545–1572.

Burns, Arthur F., and Wesley C. Mitchell (1946). *Measuring Business Cycles*. National Bureau of Economic Research.

Campbell, John, and N. Gregory Mankiw (1987). "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics*, 102, 857–880.

Christiano, Lawrence J. (1988). "Why Does Inventory Investment Fluctuate so Much?" *Journal of Monetary Economics*, 21, 247–280.

Cochrane, John (1987). "How Big is the Random Walk in GNP?" *Journal of Political Economy*, 96, 893–920.

Dickey, David A., and Wayne A. Fuller (1979). "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association*, 74(1), 427–431.

Dolmas, James (1998). "Risk Preferences and the Welfare Cost of Business Cycles." *Review of Economic Dynamics*, 1(3), 646–676.

Easterly, William (2002). *The Elusive Quest for Growth*. MIT Press.

Elliott, Graham, Thomas J. Rothenberg, and James H. Stock (1996). "Efficient Tests for an Autoregressive Unit Root." *Econometrica*, 64(4), 813–836.

Epaulard, Anne, and Aude Pommeret (2003). "Recursive Utility, Endogenous Growth, and the Welfare Cost of Volatility." *Review of Economic Dynamics*, 6(3), 672–684.

Epstein, Larry G., and Stanley E. Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4), 937–969.

Fackler, Paul L. (2005). "A MATLAB Solver for Nonlinear Rational Expectations Models." *Computational Economics*, 26, 173–181.

Fuhrer, Jeffrey C. (2000). "Habit Formation in Consumption and its Implications for Monetary Policy." *American Economic Review*, 90(3), 367–390.

Galí, Jordi, Mark Gertler, and J. David Lopez-Salido (2007). "Markups, Gaps, and the Welfare Costs of Economic Fluctuations." *Review of Economics and Statistics*, 89(1), 44–59.

Hall, Robert E. (1978). "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86(6), 971–987.

Imrohoroğlu, Ayse (1989). "The Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political Economy*, 97(6), 1364–1383.

Kocherlakota, Narayana R. (1990). "On the 'Discount' Factor in Growth Economies." *Journal of Monetary Economics*, 25, 43–47.

Krebs, Tom (2003). "Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk." *Review of Economic Dynamics*, 6, 846–868.

Krebs, Tom (2007). "Job Displacement Risk and the Costs of Business Cycles." *American Economic Review*, 97(3), 664–686.

Krusell, Per, and Anthony A. Smith, Jr. (1999). "On the Welfare Effects of Eliminating Business Cycles." *Review of Economic Dynamics*, 2(2), 245–272.

Kwiatkowski, Denis, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin (1992). "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure Are We that Economic Time Series Have a Unit Root?" *Journal of Econometrics*, 54, 159–178.

Kydland, Finn E., and Edward C. Prescott (1982). "Time to Build and Aggregate Fluctuations." *Econometrica*, 50(6), 1345–1370.

Lucas, Robert E., Jr. (1987). *Models of Business Cycles*. Basil Blackwell.

Lucas, Robert E., Jr. (2003). "Macroeconomic Priorities." *American Economic Review*, 93(1), 1–14.

McGrattan, Ellen R., and Edward C. Prescott (2003). "Average Debt and Equity Returns: Puzzling?" *American Economic Review Papers and Proceedings*, 93(2), 292–297.

Ng, Serena, and Pierre Perron (2001). "Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power." *Econometrica*, 69(6), 1519–1554.

Obstfeld, Maurice (1994). "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability." *European Economic Review*, 38, 1471–1486.

Otrok, Christopher (2001). "On Measuring the Welfare Cost of Business Cycles." *Journal of Monetary Economics*, 47, 61–92.

Pallage, Stephane, and Michel A. Robe (2003). "On the Welfare Cost of Economic Fluctuations in Developing Countries." *International Economic Review*, 44(2), 677–698.

Parker, Jonathan A., and Bruce Preston (2004). "Precautionary Saving and Consumption Fluctuations." *American Economic Review*, 95(4), 1119–1143.

Phillips, Peter C. B., and Pierre Perron (1988). "Testing for a Unit Root in Time Series Regression." *Biometrika*, 75(2), 335–346.

Phillips, Peter C. B. (1998). "Impulse Response and Forecast Error Variance Asymptotics in Nonstationary VARs." *Journal of Econometrics*, 83, 21–56.

Poterba, James M. (1998). "The Rate of Return to Corporate Capital and Factor Shares: New Estimates Using Revised National Income Accounts and Capital Stock Data." *Carnegie-Rochester Conference Series on Public Policy*, 48, 211–246.

Ramsey, Frank (1928). "A Mathematical Theory of Saving." *Economic Journal*, 38, 543–559.

Reis, Ricardo (2006). "Inattentive Consumers." *Journal of Monetary Economics*, 53(8), 1761–1800.

Romer, Christina, and David Romer (2002). "The Evolution of Economic Understanding and Postwar Stabilization Policy." In *Rethinking Stabilization Policy*. Federal Reserve Bank of Kansas City.

Stock, James H. (1991). "Confidence Intervals for the Largest Autoregressive Root in U.S. Macroeconomic Time Series." *Journal of Monetary Economics*, 28, 435–459.

Stock, James H. (1994). "Unit Roots, Structural Breaks and Trends." In *Handbook of Econometrics*, edited by R. Engle and D. McFadden. Elsevier.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2001). "The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk." *European Economic Review*, 45(7), 1311–1339.

Tallarini, Thomas D., Jr. (2000). "Risk-Sensitive Real Business Cycles." *Journal of Monetary Economics*, 45(3), 507–532.

Van Wincoop, Eric (1999). "How Big Are the Potential Welfare Gains from International Risk-Sharing?" *Journal of International Economics*, 47(1), 109–135.

Weil, Philippe (1990). "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics*, 105(1), 29–42.

Weitzman, Martin L. (2007). "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, 97(4), 1102–1130.